A Deterministic Success-Metric View:
Of Shannon's 'Weaker' Designation for the Ideal System

Paul McGough · NQP LLC
RPM@enqpy.com
Technical Note · June 2026
Companion to: Enqpy™ Formal Proof Paper (Rev 5.0, 2026) and Enqpy™ Formal Cryptographic Description (Rev 5.0) | IP Portfolio PCT/US26/22552

Abstract
Shannon's 1949 communication theory paper *("Communication Theory of Secrecy Systems")* established two secrecy constructs: Perfect Secrecy (PS), achieved by the One-Time Pad, and the Ideal System (IS), a finite-key cipher leaving the adversary permanently uncertain. Shannon designated the Ideal System *"weaker"* than Perfect Secrecy. This note examines the basis of that designation, introduces a deterministic success metric S under which both constructs produce the same deterministic S = 0 outcome, formally states the **Adversarial Outcome Equivalence Conjecture** (the general case), and demonstrates that Enqpy™ — as a proved instantiation of Shannon's Ideal System on two independent equivocation axes — achieves a structurally distinct uncertainty design from Perfect Secrecy, with two independent axes of permanent adversarial uncertainty. Shannon's "weaker" designation retains full meaning under probabilistic and cardinality-based metrics. This note argues only that it does not alter the deterministic success outcome captured by S.

1. Background: What Shannon Actually Said

In §25 of the 1949 paper, Shannon defines Perfect Secrecy as the condition in which observing any quantity of ciphertext yields zero information about the plaintext:
$H(M \mid C) = H(M)$
He then defines the Ideal System informally as a finite-key cipher that leaves the adversary with multiple consistent solutions of "reasonable probability" regardless of how much ciphertext is observed — a condition of permanent, bounded ambiguity rather than complete ignorance. He calls the Ideal System "weaker" than Perfect Secrecy and notes its analysis is "more complicated."
Shannon's "weaker" label carries two components, one mathematical and one practical:
(i) Cardinality: PS leaves the adversary facing the full set of all possible plaintexts as consistent alternatives. IS leaves them facing a proper subset — fewer alternatives, but not one.
(ii) Probability: Under a probabilistic adversarial model, a smaller set of alternatives improves the adversary's expected per-guess probability of success, even though no single guess is guaranteed correct.
This note argues that both components remain security-relevant under probabilistic and cardinality-based metrics, but are not decisive under the deliberately binary operational metric isolated here. The adversary's operational question is not "how many alternatives face me?" It is: "can I determine the answer?"

2. The Deterministic Success Metric

We introduce a binary adversarial outcome function:
Definition 2.1 — Deterministic Success Metric S

© 2026 NQP LLC. Licensed under CC-BY-4.0.*Enqpy™ · Security, Settled.*

Let A be any adversary (computationally unbounded, including quantum). Let cipher be a cryptographic primitive. Let O be the adversary's cryptographic observation record, including ciphertexts and public records such as nonces. Define:
S(A, cipher, O) = 1 **if and only if** A can uniquely determine the plaintext M from O with certainty.
S(A, cipher, O) = 0 otherwise.
S = 0 means the adversary cannot determine the answer, regardless of how many alternatives remain.
Metric S captures a deliberately binary operational outcome: whether the adversary succeeds with certainty. The size of the residual uncertainty set affects the adversary's strategy and expected effort, but it does not affect S = 0: a smaller uncertainty set does not make S = 1. S = 0 means the adversary cannot determine the answer. That is the outcome isolated by the present metric. Metric S does not replace Shannon's probability-based comparison; it isolates one deterministic operational question.
We now evaluate PS and Enqpy's IS instantiation under S.
A critical clarification before proceeding: S = 0 is not inherently strong. It has two forms. **Trivial S = 0** occurs when equivocation bounds are near-zero: the adversary cannot formally claim certainty, but one plaintext candidate sits at near-unit posterior probability. Such a system is operationally broken even though S = 0. **Substantive S = 0** occurs when the impossibility of unique determination is accompanied by independently proved large residual support. Enqpy's proved message-axis result is closed: the 2^{128} support floor holds with a uniform posterior over the full consistent set, so H = H∞ ≥ 128 bits — vast, permanent, and probability-flat. The claim in Theorem 3.1 is the conjunction: S = 0 *and* |S(CT,OR)| ≥ 2^{128} *and* H(EK,QK|T∞) ≥ 2 bits (Canonical Configuration). Those bounds are independently proved results. A near-unit posterior would make S = 0 operationally weak even without a unique solution; that gap is closed — the posterior over the consistent set is proved uniform, so no candidate sits at near-unit probability. The support floor proves the narrower claim required for S = 0: no unique plaintext solution is forced by the cryptographic observations — at least 2^{128} candidates remain consistent with any observation.

3. Proved Equivalence: Enqpy™ IS and Perfect Secrecy

The following theorem follows directly from Enqpy's existing formal proof suite (formal proof paper, Rev 5.0, 2026).
Theorem 3.1 — Outcome Identity Under Metric S (Enqpy™ Instantiation)
Model condition. Adversary A's information consists of the stated cryptographic observations and public records; the message space is non-degenerate; and external side information that determines M independently of the cipher is excluded. (Absent these conditions, S = 1 for PS and Enqpy alike before any cipher is applied.)
For Perfect Secrecy, as achieved by a correctly used One-Time Pad with uniform one-time key material:
S(A, PS, O) = 0 for all A, all O. [Shannon 1949, Theorem 1]
For Enqpy™ (proved Ideal System instantiation):
S(A, Enqpy, O) = 0 for all A (including quantum), all O.
Therefore: S(A, PS, O) = S(A, Enqpy, O) = 0 for all A, all O.
Enqpy™ and Perfect Secrecy share the same S = 0 outcome under metric S, given the stated equivocation bounds.

Proof

The right-hand clause rests on two proved results from the Enqpy formal proof suite:

© 2026 NQP LLC. Licensed under CC-BY-4.0.

Key Equivocation (KE): $H(EK, QK \mid T^{\infty}) \geq 2$ bits (Canonical Configuration), hence > 0. Within the Canonical Configuration as modeled in the Enqpy proof suite, transcript observations are proved not to collapse the admissible raw master-key origin to a unique pair (EK, QK), regardless of the quantity of ciphertext observed or the adversary's computational resources. The exact four-key value (2 bits) is the tightness that holds under known-plaintext-strength observation; ciphertext-only, the result is a permanent floor, and only positive equivocation is needed here. The key-indeterminacy result serves as supporting structure within the model. The direct basis for S = 0 is the plaintext equivocation result on the message axis.
Plaintext Equivocation (PTE): $|S(CT, OR)| \geq 2^{128}$. The adversary, given any ciphertext CT and the full output record OR, faces at least 2^{128} consistent plaintext candidates. No single candidate can be identified as uniquely correct. S = 0 follows independently on the message axis.
The message-axis support floor directly gives S = 0 for plaintext determination; the key-axis theorem gives a second, independent non-uniqueness result concerning master-key origin. The adversary cannot determine M. This is identical to the outcome under Perfect Secrecy. ■
The role of posterior probability in this analysis deserves explicit treatment. Shannon's framework tracks how posterior probabilities evolve as the adversary observes more ciphertext. That tracking is meaningful when posterior probability can, in principle, converge toward certainty with sufficient data. Enqpy's support floor permanently prohibits convergence to a unique solution: $|S(CT,OR)| \geq 2^{128}$ is an unconditional lower bound that holds regardless of the quantity of ciphertext observed, the adversary's computational resources, or the passage of time. When convergence to certainty is permanently prohibited, posterior tracking cannot change the value of S. Posterior concentration matters under finer-grained security metrics. That question is closed: the posterior over the full consistent set is proved uniform, so $H = H^{\infty} \geq 128$ bits. Within metric S, the support floor alone prevents unique determination; within Enqpy's stronger result, the uniform posterior also rules out near-unit concentration.
The proved claims are unconditional within the stated model: they do not rely on computational hardness assumptions, the difficulty of factoring, the intractability of discrete logarithms, or any assumption that may be invalidated by future algorithmic or hardware advances, including quantum computation.

4. The General Conjecture

Theorem 3.1 establishes that Enqpy's specific instantiation produces the same S = 0 outcome as PS under the stated metric and bounds. The broader claim — that any correctly constructed IS satisfying Shannon's defining conditions produces the same S = 0 outcome — is stated as a conjecture pending general proof:
Conjecture 4.1 — Adversarial Outcome Equivalence (General Form)
Let IS be any cipher satisfying Shannon's Ideal System conditions:
(a) $H(K \mid T^{\infty}) > 0$ (key equivocation is permanently positive), and
(b) $H(M \mid T^{\infty}) > 0$ (message equivocation is permanently positive; equivalently, the posterior does not collapse to a point mass on a single plaintext).
Then S(A, IS, O) = 0 for all adversaries A and all observation records O.
Therefore S(A, PS, O) = S(A, IS, O) = 0: any such IS
produces the same bare S = 0 outcome as PS under metric S. Whether that S = 0 outcome is operationally substantive depends on the size and distribution of the residual support. Shannon's 'weaker' designation retains meaning under probabilistic and cardinality-based metrics.
It does not alter the deterministic S = 0 outcome captured by metric S.
Conjecture 4.1 is stated in the general form to encompass any future cipher claiming Ideal System properties. Its proof for the Enqpy instantiation is complete under the model conditions of Theorem 3.1. The general case is expected to follow from the same logic, though a complete

© 2026 NQP LLC. Licensed under CC-BY-4.0.

general proof is deferred: in each Canonical Configuration satisfying both conditions, the plaintext equivocation bound is the direct basis for S = 0, with key equivocation providing independent supporting structure.

5. Structural Observation: The Duality Advantage

Shannon's framing implicitly treats IS as a degraded version of PS — a system that achieves "less" uncertainty. This framing inverts the correct analysis when evaluated structurally.
In Perfect Secrecy as implemented by the One-Time Pad:
Given ciphertext C and any candidate plaintext M, the corresponding key K* is uniquely determined by K* = M* ⊕C.*
This means PS operates on a single effective axis. Key uncertainty and plaintext uncertainty are not structurally independent — one is fully derivable from the other given C. The adversary's uncertainty, while complete on the message dimension, is actually constrained by the algebraic relationship K = M ⊕C.
In Enqpy's IS instantiation:
$H(EK, QK \mid T^{\infty}) \geq 2$ bits (Canonical Configuration) and $|S(CT, OR)| \geq 2^{128}$ hold independently. Neither is derivable from the other. In Enqpy, the key-axis and plaintext-axis non-uniqueness results are not merely the same uncertainty expressed twice; they arise from separate algebraic mechanisms.
Structural Observation 5.1 — Uncertainty Duality
Enqpy™ achieves two structurally independent axes of permanent adversarial uncertainty:
Axis 1 (Key): $H(EK, QK \mid T^{\infty}) \geq 2$ bits (Canonical Config) — permanent key indeterminacy
Axis 2 (Message): $|S(CT, OR)| \geq 2^{128}$ — permanent plaintext indeterminacy
These axes are algebraically distinct: the key-axis non-uniqueness and the message-axis support floor are proved by separate structures, and quotienting or resolving one axis as an interpretive matter does not eliminate the formal result on the other.
Perfect Secrecy achieves complete message uncertainty, but its key and message alternatives are algebraically coupled: for each candidate plaintext M, the corresponding key K* = M* ⊕C is uniquely determined, so choosing one determines the other. OTP therefore does not maintain an independent key-equivocation axis separate from plaintext ambiguity.*
In terms of structural design, Enqpy™'s IS is structurally distinct from PS. The standard information-theoretic sense in which PS is "stronger" concerns the completeness of the message uncertainty set. Enqpy's IS maintains a large proved message-ambiguity floor via the $|S(CT,OR)| \geq 2^{128}$ support floor, without claiming to match OTP on Shannon's original scale. The axis independence — the fact that key indeterminacy and plaintext indeterminacy are algebraically independent — is an additional structural property that PS does not realize. No dominance claim is made; the observation is descriptive.
This structural difference suggests a different response to partial compromise than OTP: PS's single-axis algebraic coupling means partial key information directly reduces message uncertainty, while Enqpy's two-axis structure does not exhibit the same one-to-one OTP coupling; formal leakage analysis is outside the scope of this note; the observation is structural and descriptive.

6. On the Probabilistic Objection

The expected objection to Theorem 3.1 takes the following form: metric S is too coarse because it collapses a cipher leaving one plaintext at 99.9999% posterior probability into the same category as Perfect Secrecy — both score S = 0, but they are clearly not equivalent in any

© 2026 NQP LLC. Licensed under CC-BY-4.0.*Enqpy™ · Security, Settled.*

meaningful security sense. This objection is correct as stated. It does not apply to Enqpy in the trivial near-zero-equivocation form described above, because the note relies on the additional stated lower bounds.

The 99.9999% scenario concerns posterior confidence. What metric S requires is identification with certainty — a unique consistent solution — and Enqpy's proof establishes unconditionally that no observation ever yields one: at least 2^{128} plaintexts remain mathematically consistent ($|S(CT,OR)| \geq 2^{128}$). For Enqpy this is sharper still: the posterior over those candidates is proved uniform (Theorem 3, closed), so no candidate approaches the 99.9999% figure the objection imagines — the near-unit-posterior case the objection relies on is provably absent, not merely deferred. These are not the same category. The objection attacks metric S as a standalone construct and is correct to do so. But Theorem 3.1 does not deploy S as a standalone construct. It deploys S as the operational question, answered by the conjunction of S = 0 and the independently proved equivocation bounds. Those bounds are what make the answer substantive.

A second form of the objection states: Shannon's "weaker" designation is not refuted by this analysis — it merely disappears under a metric chosen to make it disappear. This is philosophically accurate and conceded. Shannon's framework asks: how much does the adversary's probability of success improve per unit of observed ciphertext? Under that question, IS is weaker than PS because the residual set is smaller. The point of this note is not to refute that answer within Shannon's framework. This note isolates a different operational question: can the adversary succeed? Shannon's framework measures the rate of approach to success. Enqpy's equivocation bounds establish that certainty is permanently unreachable within the stated model. When the destination is permanently unreachable within the model, the rate of approach toward it has no consequence for metric S.

7. Scope and Consequence

The following consequences follow from Theorem 3.1, Conjecture 4.1, and Observation 5.1 within the scope of metric S and the stated model:

(i) A finite-key cipher can produce the same S = 0 outcome as Perfect Secrecy. The requirement is not an infinite key — it is permanent positive equivocation on both axes, with independently proved bounds of sufficient magnitude.

(ii) Such a cipher may realize a structurally distinct uncertainty decomposition from Perfect Secrecy, with independent ambiguity claims on key and plaintext axes that PS does not separately maintain.

(iii) The proved claims are unconditional within the stated model — they do not rest on computational assumptions and cannot be invalidated by advances in algorithms or hardware.

(iv) The construction is finite-key, reusable, and implementable in hardware and software; implementation characteristics are documented separately.

The practical significance of these consequences is developed in the companion Manifesto document. This technical note does not introduce branding or commercial terminology; it confines itself to the formal claims stated above.

8. Summary of Formal Claims

This note establishes three formal claims with different epistemic statuses:

Theorem 3.1 (proved): Enqpy™ and Perfect Secrecy share the same S = 0 outcome under metric S, given the stated equivocation bounds. This follows from the existing Enqpy formal proof suite and requires no new proof beyond citation.

© 2026 NQP LLC. Licensed under CC-BY-4.0.

Conjecture 4.1 (stated, not yet proved in full generality): Any IS satisfying Shannon's dual equivocation conditions with bounds of sufficient magnitude produces the same $S = 0$ outcome as PS under metric S. Expected to follow from the same logic as Theorem 3.1; general proof deferred.

Observation 5.1 (structural, requiring no additional proof): Enqpy™'s two independent uncertainty axes are structurally distinct from OTP/PS. Shannon's "weaker" designation retains meaning under probabilistic and cardinality-based metrics; within metric S and the stated bounds, it does not alter the adversarial outcome.

The Manifesto companion document ("The End of the Search") presents these conclusions in non-technical form for general audiences. Technical readers are directed to this note and to the Enqpy™ formal proof paper, "Ciphertext-Only Plaintext Equivocation in a Finite-Key MOD16 Record Cipher, with Exact Known-Plaintext Bounds" (Rev 5.0, 2026), which establishes the key-axis equivocation floor (Theorem 2) in full and the closed plaintext-axis result (Theorem 3): the 2^{128} support floor with a uniform posterior over the full consistent set, so $H = H_\infty \geq 128$ bits. Both axes are closed in full, with supporting derivations in Section 11 and Appendices B and C.

Enqpy™ · Security, Settled. · Proved. Not assumed.

NQP LLC · Paul McGough · RPM@enqpy.com · June 2026

IP Portfolio: P1-P6 provisionals filed and assigned · PCT/US26/22552

© 2026 NQP LLC. Licensed under CC-BY-4.0.*Enqpy™ · Security, Settled.*